

YOUR FIRST CLAUDE MANAGED AGENT

A step-by-step setup blueprint for building and deploying your first AI agent on Claude Managed Agents — the infrastructure layer Anthropic shipped in April 2026.

● MEDIUM COMPLEXITY

WHAT YOU'RE BUILDING

A Claude Managed Agent is an AI worker that lives on Anthropic's servers, wakes up when triggered, completes multi-step tasks autonomously, stores credentials securely in a vault, and reports back when done.

- Lives on Anthropic's infrastructure — not your laptop or a VPS
- Triggered by API call, webhook, Zapier, Make, or n8n
- Credential vault for all API keys — no plaintext secrets
- \$0.08 per runtime hour — most tasks cost fractions of a cent
- Notion, Asana, Rakuten and Sentry already running production deployments

BEFORE YOU START — CHECKLIST

- **Anthropic API account**
console.anthropic.com
- **API key**
Access to Claude Sonnet or Opus model
- **A defined task**
Know exactly what you want the agent to do end-to-end
- **Tool list**
Which APIs/services will the agent call?

■ **30 minutes**

That's all you need for a first deployment

STE
P 1

Define Your Agent's Job

Answer these three questions before writing any code.

- TASK** What does the agent do when triggered? "When a lead submits the contact form, research their company, draft a personalised intro email, and add it to the CRM as a draft."
- TOOLS** What APIs/services does it need? Example: Web search, Google Sheets API, Gmail (draft only)
- GUARD RAILS** What is it NOT allowed to do? Example: "Never send emails directly. Never access data outside the leads spreadsheet."

STE
P 2

Create Your Credential Vault

Managed Agents stores your API keys in named vaults — no plaintext secrets in prompts. In the Anthropic Console, navigate to [Managed Agents > Credential Vaults](#).

1. Create a new vault — name it after the project or client
Example: "lead-research-prod"
2. Add each key your agent will use:
GOOGLE_SHEETS_KEY: [your key]
GMAIL_API_KEY: [your key]
CRM_WEBHOOK_URL: [your URL]
3. Save and copy the vault_id — you will need it in Step 3.

Recommended pattern: one vault per client, or one per environment (staging vs production).

STE
P 3

Create the Agent via API

POST to the Managed Agents API with the beta header. The API returns an [agent_id](#) — save this.

```
POST https://api.anthropic.com/v1/managed-agents

Headers:

anthropic-version: 2024-01-01
anthropic-beta: managed-agents-2026-04-01
x-api-key: YOUR_API_KEY
Content-Type: application/json

Body:

{
  "name": "lead-research-agent",
  "model": "claude-sonnet-4-5",
  "system_prompt": "[Paste your system prompt – see Step 4]",
  "tools": ["web_search", "google_sheets", "gmail"],
  "vault_id": "YOUR_VAULT_ID",
  "guardrails": {
    "max_runtime_minutes": 10,
    "require_human_review": true,
    "allowed_domains": ["google.com", "linkedin.com"]
  }
}
```

STE
P 4

Write Your System Prompt

This is the most important step. Your system prompt defines the agent's workflow, tools and constraints. Use this template as your starting point:

```
You are a [role] for [Company Name].

YOUR JOB:

When triggered, you will:

1. [First task – e.g. search for the lead's company online]
2. [Second task – e.g. find relevant context about their size/industry]
3. [Third task – e.g. draft a short, personalised intro email]
4. [Final task – e.g. log everything to the CRM sheet]

TOOLS YOU CAN USE:
```

```

- web_search: research companies and contacts
- google_sheets: read lead data, update CRM rows
- gmail: draft emails ONLY — never send directly

RULES:
- Never send emails. Set status "DRAFT_READY" when done.
- Never access data outside the leads spreadsheet.
- If you cannot verify a detail, flag: status "NEEDS_REVIEW"
- Complete the full workflow before responding.

OUTPUT FORMAT (always return this JSON):
{
  "lead_name": "",
  "company": "",
  "research_summary": "",
  "email_draft": "",
  "status": "DRAFT_READY | NEEDS_REVIEW | ERROR",
  "notes": ""
}

```

STE
P 5

Trigger a Run

POST to the runs endpoint with your input data. The agent wakes up, works through the system prompt step by step, and returns the JSON output. Cost: **\$0.08/hour of runtime** — most tasks complete in 2-4 minutes.

```

POST https://api.anthropic.com/v1/managed-agents/{agent_id}/runs

Body:
{
  "input": {
    "lead_name": "Sarah Chen",
    "company": "Acme Corp",
    "email": "sarah@acme.com",
    "message": "Interested in CRM automation for our 200-person team"
  }
}

```

STE
P 6

Connect a Trigger

Managed Agents does not have native webhooks in the current beta. Choose one of these options to fire the agent automatically:

Zapier / Make
Easiest

Create an HTTP POST action that fires when a form is submitted or a spreadsheet row is added. No code required.

Replit
Lightweight

Host a tiny Express.js server on Replit that listens for webhooks and calls the Managed Agents API. Free tier handles light traffic.

n8n
Most flexible

Use the HTTP Request node to trigger agents from any workflow. Best for complex multi-step automations.

STE
P 7

Monitor and Iterate

In the Anthropic Console under **Managed Agents > Runs**, every execution is logged with timestamps, inputs, outputs and a full step-by-step trace. Filter by: COMPLETED | FAILED | RUNNING.

■ **Agent times out**

Reduce scope or raise max_runtime_minutes in guardrails

■ **Tool call fails**

Check credential vault — keys may have expired

■ **Wrong output format**

Tighten the OUTPUT FORMAT section of system prompt

■ **Hallucinated data**

Add "if unsure, return NEEDS_REVIEW" to your rules

STARTER AGENT IDEAS

● SIMPLE

■

Lead Enrichment

Name + email in, company research summary out. 1-2 tools. Great first build.

■

Invoice Classifier

Reads incoming PDF invoices, tags them by category in a spreadsheet.

● MEDIUM

- **Support Triage**
Reads inbound emails, categorises, drafts replies, flags urgent ones for human review.
- **Weekly Report**
Pulls CRM data, summarises pipeline, formats a Slack message for Monday morning.

● **ADVANCED**

- **Full Outreach Pipeline**
Finds prospects, researches them, drafts personalised emails, logs to CRM.
- **Competitor Monitor**
Scans competitor sites weekly, summarises changes, updates a tracker sheet.

Rather have us build this for you?

This is exactly what OrcaScale does for SMEs — deployed in 3-7 days, built before you pay a penny.

■ **Email us**

■ **Book a call — orcascale.co**

OrcaScale · orcascale.co · Scale operations without scaling chaos